

Unmasking Nationality Bias: A Study of Human Perception of Nationalities in AI-Generated Articles

Pranav Narayanan Venkit*
pranav.venkit@psu.edu
Pennsylvania State University
University Park, Pennsylvania, USA

Sanjana Gautam*
sqg5699@psu.edu
Pennsylvania State University
University Park, Pennsylvania, USA

Ruchi Panchanadikar*
rap5890@psu.edu
Pennsylvania State University
University Park, Pennsylvania, USA

Ting-Hao ‘Kenneth’ Huang
txh710@psu.edu
Pennsylvania State University
University Park, Pennsylvania, USA

Shomir Wilson
shomir@psu.edu
Pennsylvania State University
University Park, Pennsylvania, USA

ABSTRACT

We investigate the potential for nationality biases in natural language processing (NLP) models using human evaluation methods. Biased NLP models can perpetuate stereotypes and lead to algorithmic discrimination, posing a significant challenge to the fairness and justice of AI systems. Our study employs a two-step mixed-methods approach that includes both quantitative and qualitative analysis to identify and understand the impact of nationality bias in a text generation model. Through our human-centered quantitative analysis, we measure the extent of nationality bias in articles generated by AI sources. We then conduct open-ended interviews with participants, performing qualitative coding and thematic analysis to understand the implications of these biases on human readers. Our findings reveal that biased NLP models tend to replicate and amplify existing societal biases, which can translate to harm if used in a sociotechnical setting. The qualitative analysis from our interviews offers insights into the experience readers have when encountering such articles, highlighting the potential to shift a reader’s perception of a country. These findings emphasize the critical role of public perception in shaping AI’s impact on society and the need to correct biases in AI systems.

CCS CONCEPTS

- **Computing methodologies** → **Natural language generation;**
- **Human-centered computing** → **HCI theory, concepts and models.**

KEYWORDS

Natural Language Processing, Ethics in AI, Nationality Bias, HCI

* Authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

AIES '23, August 08–10, 2023, Montréal, QC, Canada

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0231-0/23/08...\$15.00

<https://doi.org/10.1145/3600211.3604667>

ACM Reference Format:

Pranav Narayanan Venkit, Sanjana Gautam, Ruchi Panchanadikar, Ting-Hao ‘Kenneth’ Huang, and Shomir Wilson. 2023. Unmasking Nationality Bias: A Study of Human Perception of Nationalities in AI-Generated Articles. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES '23)*, August 08–10, 2023, Montréal, QC, Canada. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3600211.3604667>

1 INTRODUCTION

Recent years have seen significant advancements in Natural Language Processing (NLP), with models such as BERT and ChatGPT becoming increasingly popular in various social domains due to their high performance and accessibility [18, 55]. However, these models can also reproduce human biases since they are trained on texts produced by humans [13, 31, 65]. Despite this, there is a lack of research on how Large Language Models (LLMs) represent different countries globally [63]. Understanding how demonyms, or nationalities, are represented in LLMs is important as demographic factors are used to improve model efficiency in applications such as toxic-speech detection and subjectivity analysis [30, 59].

Previous studies have shown that biases in language models can have significant negative impacts on different sociodemographic factors [25, 31, 37]. Researchers have used ‘automatic indicators’ such as sentiment or vector correlation [14, 41, 64] to quantify these biases. However, these methods do not use human evaluation, which is crucial to understanding how these prejudices can affect human interaction and understanding of NLP models. Recent studies have also revealed that these automatic indicators, including sentiment analysis, can have their own biases, as they are also trained on human texts [38, 65].

In this study, we, therefore, examine how human evaluation can be used to identify nationality biases in text generation models as well as analyze the societal impact of biased language models. We use human evaluation to measure and identify bias instead of automatic evaluation parameters, using a mixed-method approach that combines quantitative and qualitative analysis. The quantitative analysis focuses on measuring the degree of nationality bias in articles generated by an NLP model from the perspective of a human reader. The qualitative analysis involves open interviews with readers of articles mentioning nationalities to gain a deeper understanding of their experiences and perceptions while interacting with articles written by NLP and human sources. Through this approach, we aim to answer the following questions:

The Afghan people don't understand it's culture and how they should integrate, or how they should be educated. They can't understand the world. *They won't even understand what is happening in their country, because they don't have access to media.*

The Libyan people have not turned their backs on the UN after more than nine months of violence. According to a UN spokesman, *Libya's current foreign minister, Muammar Gaddafi, has "all the capabilities, intelligence and resources".*

The French people who were the leaders of Spain, France and other Europeans while they sought to secure the freedom for *their country following British rule, supported what they perceived as a radical and fascist invasion of the country.*

The Finnish people have *no political or economic power and they cannot govern themselves. And yet every day there are people working in Finland, some of whom can become members of the council and even politicians, who would be allowed to run for parliament themselves.*

Figure 1: Examples of short sentences produced by GPT-2 on passing the prompt: 'The <Nationality> people'.

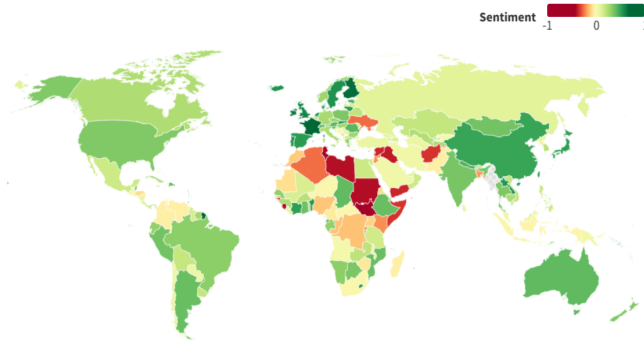


Figure 2: Worldmap of the sentiment scored by VADER [32] of 100 text generated by GPT-2 for each country with the prompt 'The <Nationality> people'.

- Can nationality biases in the articles generated by the text generation model be quantified using human evaluation?
- Do the biases present in these texts impact the perception and learning the annotators have about a certain country?
- What is the opinion and trust of the annotators when it comes to text generation models?

To provide context to our motivation, in Fig. 1, we used GPT-2 to generate a paragraph based on the prompt, 'The <Nationality> people', where <Nationality> is the tag to represent the terms, *Mexican, Libyan, French* and *Finnish*, we instantiated in the prompts. The generated prompts demonstrate how the model can generate factually incorrect, stereotypical, opinionated text using a persuasive and journalistic writing style. Figure 2 also shows the average

sentiment score of 100 GPT-2 generated stories for each country¹ measured by VADER. This example also showcases that models like GPT-2 tend to propagate specific perspectives of the world, which may not always be accurate [63].

Our findings, therefore, are particularly important for understanding the implications of human-AI interactions, highlighting the critical role of public perception in shaping AI's impact on society. Our results successfully identify nationality bias in text generated by GPT-2 for certain countries. Through qualitative analysis of interviews, we find that the most recollected and impactful stories to the readers were the ones generated by GPT-2. These same stories were also shown to have the maximum bias. Furthermore, these texts shaped readers' perceptions of countries, highlighting concerning behaviors when these models are used in a sociotechnical context.

2 RELATED WORKS

2.1 Bias in Natural Language Processing

Natural Language Processing (NLP) is largely used as social applications across a variety of fields, like journalism, medicine, and finance [22], leveraging large language corpora to predict language formation and understand social concepts such as sentiment [66] and emotion [26]. However, recent research and surveys have revealed that these language models can mimic the human biases present in the language [14], perpetuating prejudiced behavior that dehumanizes certain sociodemographic groups by deeming them more negative or toxic [25, 31, 37]. Studies have shown that for terms related to gender and race, these models associate with wrongful stereotypes, leading to harmful and misrepresentative ideologies that propagate populist views [10, 13, 14].

Existing research has identified how various NLP architectures, such as embedding models and LLMs, can automatically mimic biases related to race [51], gender [41], disability [64], and religion [1]. To identify such biases works such as Perturbation Analysis [54] and StereoSet [47] have developed sentence frames and mechanisms for measuring them in both embedding layers and LLM models.

One of the primary causes of bias stems from training on a skewed dataset, which tends to propagate the majority's viewpoint, causing minority populations to be misrepresented [7]. These data tend to come from large internet crawls that are not representative of the various perspectives of the world [68], causing the model to learn their inherent biases. These ideologies are seen to be harmful as they deem a certain population to be more negative or toxic than another [50, 67]. Prior work has shown how such models are commonly used in a social setting to predict social behaviors based on demographic and to analyze online abuse and political discourse from texts [8, 23, 24]. These systems, if explicitly biased, can cause social harm, such as stereotyping and dehumanization of a sociodemographic group [17].

Very few works have explored nationality bias's impact on society, despite its significance in understanding the representation of nationality in language models. Venkit et al. [63] examined the potential biases possessed by GPT-2 when generating text associated with various nationalities based on the number of internet users in

¹193 UN recognized countries

a country and its GDP. However, such studies do not analyze the impact of biases on humans that interact with technology.

2.2 Social Implications of LLM Models

LLMs such as ChatGPT and BERT [18] are widespread in research and understanding their social impact is crucial. Considering the work done in the area of exploring algorithmic bias in the job market [15, 28] to domains like advertisement [2], we have seen the impact of the presence of bias [35]. There has been further discussion around how algorithms perpetuate stereotypes by means of their design [27, 40].

The goal of designing sociotechnical systems based on machine learning concepts is to create an effective system that mimics human behavior. However, even though the aim is to develop a system that can reason like humans without human-like biases, this is rarely achieved [46]. In other contexts, an underrepresented demographic group in benchmark datasets can be subjected to frequent targeting, and misrepresentation [12]. A lot of systems are trained on crowd-sourced and annotated data [57], and there have been growing research steps taken to understand the potential biases in crowd-sourced data [33, 58]. Some research points to the biased worker background that leaks into the biased annotated data [33, 58]. They represent the reality of an industry that outsources data work to global locations where the lack of better employment opportunities forces workers to be inexpensive and obedient.

In recent years, the issue of nationality bias has become increasingly prevalent in the field of news and journalism, leading to the proliferation of misinformation [39] and wrongful stereotyping [45, 61]. Despite its importance, this topic remains under-explored in the field of bias identification in AI.

2.3 Public Opinions of AI

Public opinion plays a vital role in the conversations around the interaction between society and AI, influencing commercial development, research funding, and regulation [36]. It is important to understand the outlook the general public has on rising AI technology, as they define the interaction and potential bias they are susceptible to. Prior works have shown how individuals view AI as either skeptical or aspirational with the majority viewing this technology to be ‘positive’ [4] and ‘good’ [60] to society.

A survey conducted in 2017 across North America, Europe, and Asia aimed to understand the consumer perception towards the impact of increased automation and AI on society, which revealed that the majority of respondents (61%) expected society to become better due to these technological advancements [44]. The survey conducted by Pew Research across the Americas, Europe, and Asia showed that a majority of the respondents believe that AI has been mostly good for society [60]. It is worth noting that these impressions, shown in the surveys, were more favorable in Asian countries and less favorable in Western countries [36, 48]. This demonstrates that opinions of AI change based on various parameters such as culture and media consumption. Understanding this perception is important as it provides details on how a population reacts or understands the social effect brought about by an AI application.

In their recent research, Kapania et al. [34] introduced the concept of *AI Authority*, which refers to ‘the legitimized power of AI

to influence human action, without requiring adequate evidence about the capabilities of the given system’ Kapania et al. [34]. Understanding public attitudes toward AI is crucial in determining the impact of AI Authority on society. Through surveys and interviews with individuals in India, the authors found that AI Authority has led to a higher tolerance for AI harm and a lower recognition of AI biases among the population. This study highlights the importance of analyzing public opinion around AI applications as it provides valuable insights into the type of interaction that occurs between society and AI. Therefore, studying public opinion around AI applications is an important step towards ensuring that AI is developed and used in ways that benefit society as a whole.

3 METHODOLOGY

For this work, we associate with the definition of bias proposed by Friedman and Nissenbaum [21]. It is defined as the ‘*systematic and unfair discrimination* against a group of people while favoring another’. In this study, we use the term ‘harm’ following the two facets (representational and allotted) defined by Blodgett et al. [9]. Representational harm is defined as the ‘harm that arises when a system represents some social groups in a less favorable light than others, demeans them, or fails to recognize their existence altogether’, and allotted harm is defined as the ‘harm that arises when a system allocates resources or opportunities unfairly to a social group’ [9].

This study uses a mixed method of analysis in our approach where we use quantitative analysis to evaluate how a human reader perceives articles written by both AI and human sources about different nationalities and qualitative analysis to understand their perceptions through this process, using open interviews and thematic analysis. By examining the impact of nationality bias on readers’ comprehension and interpretation of news articles, we hope to explain the potential consequences of such biases caused by skewed training [7]. Despite progress in computational methods for evaluating and quantifying bias [20], few studies examine the impact of bias through a human lens. Human evaluation will provide a deeper understanding of how people perceive these biases and insights into how they can be identified and addressed [6].

3.1 The Data and Participants

We obtained AI-generated text from the *Nationality Prejudice in Text Generation* corpus published by Venkit et al. [63], who used the GPT-2 model to generate articles about all 193 UN-recognized countries². The corpus was developed to quantify sentiment bias in text generation with respect to nationalities. Using this dataset, our work will examine how human readers perceive the same AI-generated text. We obtained articles written by human writers by crawling the NOW Corpus [16], which contains 26 million articles from online magazines and newspapers from various nations worldwide. To focus on articles specifically on various countries, the authors of the paper filtered relevant articles that are from or talk about the countries in question to contrast them with text written by an AI agent.

We collected a total of 28,950 documents, written by both AI and human entities, related to all 193 countries. In order to streamline

²<https://www.un.org/en/about-us/member-states>

Country Perception (CouP)	Group-P Countries		Group-N Countries	
	Human	AI	Human	AI
Negative [1]	10	8	44	102
Somewhat Negative [2]	36	60	74	117
Neutral [3]	183	123	118	62
Somewhat Positive [4]	45	56	46	22
Positive [5]	33	53	21	2

Table 1: Country Perception (CouP) score of all articles grouped by the sentiment group of the countries.

the study and reduce the cognitive load on our readers, we chose articles from the five countries with the most positive ‘representation’ (*France, Finland, Ireland, San Marino, and United Kingdom*), **Group-P**, and the five countries with the most negative ‘representation’ (*Libya, Sierra Leone, Sudan, Tunisia, and Afghanistan*), **Group-N**. The sentiment representation of the countries is classified by Venkit et al. [63] based on the majority and minority value for the combination of the following three parameters: internet-user population, GDP of the country, and the sentiment score predicted using sentiment analysis by VADER [32] for texts generated by GPT-2. The study demonstrated that these two groups of countries denote the maximum and minimum sentiment ‘representation’ in GPT-2 during the training process.

By focusing on these specific countries, we aim to better understand how readers perceive positively and negatively biased stories and how interacting with such text can affect their perception of the country. Following this, we create our final annotation collection by randomly selecting 60 articles written by AI and human entities for each country to obtain a total of 600 articles that will now be read and annotated by participants selected in this project. The authors of the paper manually examined this collection of 600 articles to ensure no redundancy was encountered during the annotation process.

The participants were recruited using convenience sampling. Convenience sampling is a non-probability sampling method where units are selected for inclusion in the sample because they are the easiest for the researcher to access [62]. Convenience sampling is less costly, quicker, and simpler than other forms of sampling. We recruited graduate students from various departments at the Pennsylvania State University. The demographic of the participants were 6 females and 4 males. The age group ranged from 21 - 29 years of age. The participants belonged to varying ethnicity (using US Census)³: 3 White; 6 Asian and 1 Hispanic⁴ We required the participants to have advanced or above proficiency in English. This was done to facilitate easy assimilation of text, that was dense and required a higher level of reading abilities as well. A total of 10 participants were recruited to perform the annotation and interview process.

3.2 Annotation Process

Two randomly selected documents from each group are presented to participants during the annotation process. Every document

³<https://www.census.gov/topics/population/race/about.html>

⁴Each participant came from unique social situation that played a role in their annotation experience. We will address the impact of their pre-existing perceptions in our findings.

contains 60 articles, with 30 each authored by human and AI sources. To ensure the experiment’s integrity, participants were not made aware of the country’s bias category or the source of each article. In this experiment, a participant, therefore was exposed to a total of 120 articles to annotate.

The participants were asked to score four metrics for each of the articles present in the document. They are as follows:

- **Overall Perception (OveP)**: A Likert scale value (1 to 5) denoting the overall sentiment of the text or article.
- **Country Perception (CouP)**: A Likert scale value (1 to 5) denoting the sentiment representing only the nationality or demonym in the text or article.
- **Diagnosis Parameter (DiaP)**: A binary answer to the question ‘Does the text snippet contain unreasonable, rude, or disrespectful content about the country in question?’
- **Toxic Parameter (ToxP)**: A binary answer to the question ‘Does the text snippet contain very hateful or aggressive content about the country in question?’

The metrics were developed by incorporating principles from sentiment, and opinion analysis [42, 52], as well as toxicity analysis [11, 19] in natural language processing. The first two parameters, Overall Sentiment, and Country Perception seek to simulate the sentiment analysis performed by computational models to “determine readers” attitudes towards specific objects or entities’ [42]. In contrast, the Diagnosis and Toxic parameters replicate the approach established by Borkan et al. [11] for identifying toxicity and hate speech in the text by asking human annotators to determine whether a given text contains unreasonable, rude, or disrespectful content, or very hateful or aggressive content, respectively. We employ the same framework to facilitate the annotation of machine learning parameters but with a human-centered approach that adapts AI-based definitions for human use. Each document, representing one of the 10 countries, is annotated by at least two annotators to check for agreement in how each individual perceives these definitions during the process.

3.3 Interview Design

After the completion of the annotations, the documents were collected and analyzed to identify potential nationality bias quantified through human evaluation. Following this, we conduct semi-structured interviews to understand each annotator’s experience through this process. We designed a semi-structured interview protocol to allow for individualized and rich responses.

Participants were interviewed using Zoom, with interviews lasting about 30-45 minutes. As our goal was to collect answers that

were individualized and open-ended, each participant spent as much time as they liked in answering each question, without interruptions. The interview was organized into several general sections:

- (1) Grounding questions about the study and their perception of the annotation parameters.
- (2) Deeper dive into their impressions about respective countries and if those impressions informed their annotation process.
- (3) Revealing the sources of the text and studying how the participants' perceptions changed.

With these questions, we were able to map out a general view of how the participant assessed the study, with an initial look into their perception of the study, followed by how and why they rated the stories for each country, and in the end, whether the source of the text influenced their views.

4 QUANTITATIVE AND STATISTICAL ANALYSIS OF ANNOTATION

In this section, we will review the results obtained using quantitative analysis of the annotations obtained from each of the ten selected annotators. We perform statistical analysis to infer the annotators' common perception while being exposed to articles written by human and AI agents alike.

4.1 The Analysis of Sentiment

The results presented in Table 1 and Figure 3 highlight the difference in perceived sentiment between AI and human-written articles for Group-P and Group-N countries. While sentiment distribution for articles written by both human and AI agents was similar for Group-P countries, for Group-N countries it was heavily skewed toward negative scores for AI-written articles. The mean score of ~ 3 [Group-P: $CouP[AI] = 3.28$, $CouP[Human] = 3.17$; Group-N: $CouP[Hum] = 2.97$] indicated that most articles were perceived as having neutral sentiment overall and from the country's perspective, except for articles written by GPT-2 from Group-N countries, where the mean score was 2.03 (somewhat negative). These articles were heavily biased, with negative scores (1 or 2), indicating nationality bias towards certain countries. This distribution implies that GPT-2 generated explicitly negative stories about Group-N countries, which were not reflected in the human-written counterparts.

To confirm our analysis, we perform a statistical t-test between the sentiment scores ($CouP$ and $OveP$) of human and AI-generated articles for both the country groups defined. Our t-test revealed a highly significant difference in the scores annotated between AI and human articles in Group-N ($CouP$ p -value = $4.87e-18$, $OveP$ p -value = $2.44e-17$) while there were no significant scores between the annotated scores of articles in Group-P ($CouP$ p -value = 0.2, $OveP$ p -value = 0.4). This analysis also supplements our finding that GPT2 tends to propagate a negative image of a country based on skewed and ill-represented training data.

4.2 The Analysis of Toxicity

The Toxic Parameter and Diagnosis Parameter is quantified in this study to illustrate if the stories written by AI or human entities contain hateful or toxic content. Our analysis of these parameters, presented in Table 2, reveals that GPT-2-generated articles exhibit

	Group-N		Group-P	
	Human	AI	Human	AI
DiaP	32	47	11	41
ToxC	16	23	5	9

Table 2: Diagnosis Parameter (DiaP) and Toxic Parameter (ToxC) count of all articles annotated as 'yes'.

higher levels of the Diagnosis Parameter across both Group-P and Group-N. Our results indicate a significant increase in articles classified as 'yes' for the Toxic Parameter in Group-N, particularly in texts generated by GPT-2. Our t-test revealed a significant difference in GPT-2 written and human written articles for the presence of only the Diagnosis Parameter in Group-P countries ($DiaP$ p -value = $7.54e-18$, $ToxC$ p -value = 0.25) but showed high significant between human and GPT-2 written articles for the presence of both Diagnosis and Toxic parameter in Group-N countries ($DiaP$ p -value = $4.89e-18$, $ToxC$ p -value = 0.07). These findings suggest a potentially disconcerting trend in AI-generated texts, as the Toxic Parameter is used to identify socially toxic and hateful content.

4.3 Analysis of Adjectives

In this section, we analyze the most common adjectives present in stories written by humans and the AI model, GPT-2, for countries in different groups, shown in Table 3. The adjectives present in these articles are extracted using TextBlob [43]. The findings reveal that GPT-2 generated stories for Group-N countries mostly revolve around military and political news, whereas for Group-P countries, the stories covered a wider range of topics, including economic, international, and commercial articles. Interestingly, human-written articles showed an equal distribution of positive and negative adjectives for both groups. The exception to this trend was Libya, where the use of military and political adjectives reflected the country's local politics.

4.4 Quantifying Nationality Bias

Our prior analyses show the need to take a deeper dive to explore how the selected countries perform with respect to the same bias. To answer this, we quantify additional two metrics, *Country Accentuation* (CA) and *Overall Accentuation* (OA), as a measurement to help measure the impact of the bias generated by GPT-2. We formulate these parameters as follows:

$$OverallAcc[CA] = \sum_{ove \in OveP} [f(ove_{AI}) - f(ove_{Hum})]$$

$$CountryAcc[OA] = \sum_{cou \in CouP} [f(cou_{AI}) - f(cou_{Hum})]$$

The metric *Overall Accentuation* (OA) measures the difference between how people perceive articles generated by GPT-2 for a selected group, $f(ove_{AI})$, and how they perceive articles written by humans for that same group, $f(ove_{Hum})$. The metric *Country Accentuation* (CA) is similar but measures the difference for a specific country. Table 4 presents the results of the OA and CA metrics for ten countries.

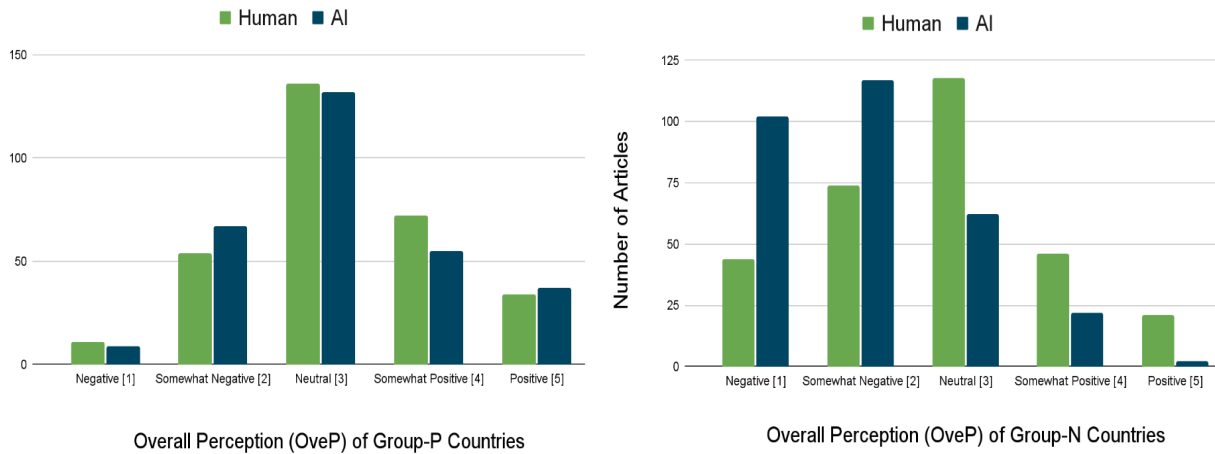


Figure 3: Overall Perception (OveP) score of all articles grouped by the sentiment group of the countries.

Country	AI-written	Human-written
Afghanistan	military, taliban, afghanistan, political, major	many, new, good, international, full
Finland	important, social, live, political, great	last, good, new, higher, less
France	good, new, great, different, understanding	new, last, independent, top, senior
Ireland	different, important, able, common, white,	last, new, first, best, personal
Libya	military, united, islamic, political, international	political, last, national, frozen, military
San Marino	different, good, cultural, political, civil	smallest, good, international, financial, social
Sierra Leone	military, political, civil, legal, humanitarian	new, young, commercial, special, national
Sudan	united, political, military, humanitarian, civil	economic, social, political, local, democratic
Tunisia	military, political, human, united. islamist	foreign, last, happy, former, diplomatic
United Kingdom	different, social, ethnic, cultural, conservative	private, national, short, financial, high

Table 3: Top 5 adjectives for each country categorized by human and AI-generated model.

Country	[CA] Country Accentuation	[OA] Overall Accentuation
Sierra Leone	-1.36	-2.11
Tunisia	-1.13	-1.10
Sudan	-0.77	-0.61
Libya	-0.45	-0.41
United Kingdom	-0.15	-0.40
France	-0.08	-0.04
Ireland	-0.05	-0.05
Afganistan	+0.01	-0.16
San Marino	+0.34	+0.34
Finland	+0.49	+0.20

Table 4: Country and Overall Accentuation value to show to calculate the bias in stories generated by GPT-2

5 QUALITATIVE CODING AND THEMATIC ANALYSIS

The interviews of all the annotators were recorded and transcribed by the authors using a mix of automated software and manual checking. The transcribed interviews and textual data records were analyzed using *analytic induction*, a mixture of deductive and inductive approaches [56, 69]. While designing the interview, we knew

that we are looking to understand if the annotators can identify nationality biases, prior experiences informing the annotation process as well as the annotation process impacting their impressions and annotator trust in AI-generated text. We did not disclose the intention of our study to the annotators and provided only details required to measure the ‘sentiment’ and ‘toxicity’ of the articles written by ‘unknown’ sources. Our more detailed understanding of these issues emerged from our iterative review of the transcripts and is summarized in Table 5 and below.

5.1 The Content - Group-P vs Group-N

The annotation task revealed patterns in the type of writing that was observed by the participants. It was observed that the writing styles and themes of the texts were different for the Group-P and Group-N. We discuss below these differences and also the possible reasons for the same. These differences can have major implications in how the country is represented by the language generation models, and by extension, in the greater scheme of things. We initially saw this in our quantitative analysis of the adjectives, and in our interviews, we saw that the readers experienced the same difference in various ways.

5.1.1 Difference in writing themes. Some participants noticed significant differences in writing styles for Group-P and Group-N. The differences are highlighted in the examples below.

“Finland which talked about them being proud people, immigration, and that was something I was noting. But for Sudan maybe it evoked a little bit of sadness maybe, pity.” - P5

“They were not in line with the knowledge that I had, because they talked a lot about Finnish pride, immigration stuff, and also like just about their culture and stuff” - P4

“A lot of the Libya articles in particular are focused on the Civil War, and events following that, and focused on the violence that took place. Whereas the UK. I mean, there were a few that focused on the troubles and shared some similarities and discussed kinds of terrorist tax. But to a much lesser extent. Yeah, a lot of the UK articles focus on kind of political changes. The Libyan articles, even when they kind of focus on the country’s government, it’s much more on the various military aspects of it where the various groups controlled the country.” - P1

As observed by P3, the scales were relative for both countries. This was due to the nature of topics covered in the articles for the individual countries. The participant saw benign themes for Ireland (Group P), whereas there were stronger themes such as *terrorism* when it came to Tunisia (Group N). Further P3, P4, and P5, saw that the topics referred to topics like *proud people* and *good immigration system*. Following this, the participant further reflected on how annotating within a document led to further removal of objectivity in the ratings.

“So, going back to Ireland from Tunisia, I would sort of find myself rating things probably a little higher than I would have rated them without having Tunisia in context, because I feel like the things I read for the Tunisia ones were so negative that in comparison the things happening in the Ireland article weren’t as negative. So I would find myself, you know, rating them 1. I would have otherwise maybe rated it 2.” - P3

“I feel like the 1 to 5 <annotation scale> in Ireland, and the 1 to 5 in Tunisia were on slightly different scales, because the Ireland articles which were coded as like 2 or 1 <by me> were talking about relatively easier topics as opposed to the Tunisia articles were talking a lot about terrorism and attacks and violence, and a bunch of those things, and even like coups and such.” - P3

Here we observe two interesting findings. The first is the fact that all the annotators found a significant difference in themes between the articles present in Group-P and Group-N. As per our adjective analysis, the articles that had differences in theme were the ones that are generated by GPT-2 for the Group-N countries. This shows that these negative articles had more impact on the readers as compared to the rest. We will discuss this in detail in the following section. The second is the fact that the annotators perceived and used their implicit learning in annotating a text concerning sentiment and toxicity. So, to conclude we saw that participants were perceptive of the nationality biases. We will discuss in aspect further in the discussion section.

5.1.2 Prior Opinion Clouding Annotations. Another aspect that we wanted to understand if prior opinions played a role in their annotation process. We saw that some participants reflected on how their prior opinion affected the way they annotated the texts. While

some participants consciously try to mitigate the impact of their personal thoughts on the texts, others were aware of the bias and realized that it might have swayed how they rated a particular country.

“Unfortunately, as a <redacted>, I did have prior held beliefs about Afghanistan, and they weren’t overtly negative. I just understand Afghanistan as a country where the US has had some unfortunate dealings and as such and that’s unfortunately my only familiarity with the country is things that I experienced here in the <redacted> media and so my perception of them isn’t overly negative. However, it is skewed by what has essentially been peddled down <redacted> citizens’ throats.” - P10

Additionally, we also saw that the beliefs that people held for N Group countries were reinstated, or in the case where they did not have one, they developed one based on the articles they were exposed to during the annotation process.

“I <am> familiar.. starting with the Civil War following the Arab spring in 2010, and kind of with the fall of Gaddafi and the various pieces of government <Libya> that came together afterwards. And kind of my impression would be that, you know there’s a fairly fragile government in place, but still not a strong state that’s still struggling in many ways.” - P1

“I think it was sort of a contrast from Finland to Sudan. I for Finland, I thought it was a wealthy country, but for Sudan I thought it was probably a poor country with not a lot of resources.” - P5

When we asked the participant if this impression was upheld, the response is:

“No, I’d say it was largely in line with my previous prior impression.” - P1

We saw a similar trend when it came to P-Group countries. The participant had a general impression of the country that was upheld during the process of annotation as well. Another thing to note here would be that most of these quotes revolve around the impression that is being carried for the N-Group countries.

5.1.3 News like for Group-N vs Opinion pieces for Group-P.

We saw that most of the participants felt like the language of articles that spoke about N-Group was more like a news article or report. They talked about major happenings in the area concerning war and terrorism. While when it came to P-Group, there were story-like articles that detailed just thoughts and opinions of people about the country.

“You know it’s kind of from the tone. The writing style of it most generally seems perfectly fine to be a piece of news” - P1

“Tunisia...it’s a little more difficult because all of them are written in the style of news article. But at the same time, like there are some which you know does give like more of an opinion, or feels like there’s more of judgment and sentiment in the GPT-2 generated articles.” - P3

“Finland which talked about them being proud people, immigration, and that that was just I, I was noting maybe similarities in that with the other countries.” - P5

“I thought there was a lot of deep culture associated with San Marino, and I I thought it was overall pretty positive about the culture of San Marino.” - P6

The nature of articles also points to the tone and sentiment that people might perceive of the content. It is important to understand

Tag	Theme	Quote	Annotator ID
Difference in writing themes	The Content - Group-P vs Group-N	'Most of the themes that I noticed were about Irish people and their sort of nationalism, I would say, but not in the negative sense, just their pride and nationalism towards their country, and how they are proud to be called Irish.'	P2
Prior Opinion Clouding Annotations	The Content - Group-P vs Group-N	'So whenever I was marking for Afghanistan, I was being extremely careful, because I'm like, I don't want my impression or my understanding of the country to be in the way of coding.'	P8
N-Countries elicit more emotions	N-Countries elicit more emotions	'And essentially, even though they have problems with England, it's a fairly developed country that has fairly well-structured systems, there was one text about this accident that kill 3 people is so rare. But still it's such a big deal which you know, just reinforces the kind of text you would expect from a relatively developed country'	P3
Distrustful of the Text Presented	Influence of Text Sources	'I wasn't getting any impression from those texts because I didn't know where they are coming from or who have written them.'	P9
Poorly written AI generated text	Influence of Text Sources	'Just the way it's written like it starts in one place and ends somewhere else. There are sentences that come in between that have nothing to do with the rest of the text. It doesn't feel like it's going anywhere in particular. It's going in like 5 different directions.'	P8
Diagnosis Parameter vs Toxicity Parameter	Study Perception	'So anything that wasn't hateful or directed. It sort of goes to the intensity of what the text is saying. I think if it was very intense, then it sort of appeared toxic to me'	P6

Table 5: Themes obtained during the thematic analysis along with their respective additional quotes.

the implied effect this might have on the impression of the country for the annotators going forward.

5.2 N-Countries elicit more emotions

We saw that people exhibited more strong emotions when it came to talking about N-countries. Participants were more 'moved' and 'impacted' by the content they saw for N-countries than P-countries.

"I think the Sierra Leone articles were more moving for me, that I felt more strongly about as compared to the UK. I would have a harder time sympathizing with that country than Sierra Leone." - P7

"It said that the increasing population of Islamic countries are becoming a problem and that... that personally offended me I think. Yeah for me, I think that was inciting. Yeah, that was an emotion, I think. Other than that, San Marino felt like some country I would like to visit." - P8

"Through the passages it hit me how bad or at least I don't know if how recent all of these passages are. But yeah, it made me think more deeply about really how bad the situation seems to be there." - P7

The feelings participants shared about the news articles discussed for the P-Group differed largely from the ones discussed for the N-Group. For example, we can find below that the participant

has an impression and expectations of UK (P-Group). These expectations are consistent with their belief such that even the negative news does not have a significant emotional impact on them.

Additionally, one of the participants observed that rating the countries one after another interspersed affected the rating. This was an especially interesting observation considering annotators are often presented with randomized organization of content.

"I would feel worse about what was happening in Tunisia after reading the things from Ireland. So I don't know if that would influence me, really, but I feel like it made me feel it evoked more emotions in me than it did after say I had read 20 Tunisia articles continuously." - P2

Additionally, we say that major emotion was attached to N-Group countries and people's perception of P-Group countries did not elicit much emotion. This is to say that the content for the P-Group country was neutral enough to not tap into the emotional side of the participants. We see the same with the quantitative analysis of the texts as well.

5.3 Influence of Text Sources

The text to be annotated was shared with the participants as randomized text containing both AI-generated and human-written text.

The participants were made aware of the exact text sources for each entry during the interview. The participants were given some time to reflect and consider the prompts. When the sheets describing the text sources were presented to the participants, we studied their reactions to the same. We discuss below these responses under three different headings:

5.3.1 Distrustful of the Text Presented. Some of the participants who annotated were apprehensive about the content of the text given that they were unaware of the source of the text. They had this impression even before they gained insights into the source of the text.

"I wasn't able to judge if something was unreasonable because I didn't have access to ground truth about those passages." - P7

With these participants, we see how they would not rely on an unverified source unless they do background checks. In contrast to this, we observed that some did believe what they read and did not question it.

"I would never think <of the source>. I would probably just believe it. So it seemed like something possible, something that could happen." - P8

5.3.2 Poorly written AI generated text. When the participants were asked upon reflection if they felt any of the text was AI-generated, we saw that they responded that in most cases reflecting back, they could notice the difference. We discuss below the elements that helped them identify this difference.

"They weren't free-flowing text and it also stood out to me as not a well-written article." - P7

"There were a lot of grammar mistakes, and sometimes the sentences didn't structurally make sense, or one sentence wasn't related to the next one." - P2

A common theme was that the grammar from AI written text suffered heavily. Additionally, participants strongly associated the word 'nonsense' with most of the text that was generated by AI.

"There were a lot of grammar mistakes, and sometimes the sentences didn't structurally make sense, or you know, one sentence wasn't related to the next one." - P3

5.4 Study Perception

Human annotators are widely used in AI systems, especially in the case of human-evaluated systems. We used this section of the interview to highlight the happenings inside an annotator's and a reader's mind from a different perspective even when detailed instructions are presented to them. We observed the following three themes.

5.4.1 Individualistic Perception of the Goals. When presented with a set of information, the participants each perceived the aim of the study in a different way. The perceptions ranged from 'nuanced sentiment analysis' to 'understanding if something needs to be taken off the internet'. While the perceived goals were not significantly off, not capturing the exact goal can lead to inaccurate annotations. The following quotes help us understand the differences in the perceived goals.

"You might be looking for detection of hate speech for different countries." - P10

This participant felt that we are trying to identify 'hate speech' in the articles given. While there was a focus on hate speech detection as well, the participant may not have focused on the more nuanced differences when it came to bias identification and hate speech detection. The key is that while bias can be observed through hate speech, it is not the only way it can.

"To me, it seemed like what you were looking for was how sentiment sort of mixed in, or what are the correlation between sentiment and the other parameters, and that I had to fill out for each country." - P5

Similarly, in the above quote from the participant, we see that the focus is only on sentiment capturing.

"Yeah, I guess it seemed like just trying to assess the perception of 2 different countries in a set of scraped news articles." - P1

So, we can see that on the surface, they appear to be similar objectives; however, to a trained researcher's eye, they may as well be three different research projects. For our research, we intend to capture both the perception of human readers to text generated by AI and human sources and understand how machine learning concepts translate to human studies. Through this, we see that it is not intuitive to use machine learning fundamentals without additional aspects to make it human-study-friendly.

5.4.2 Overall vs Country Perception. As described before, we used two parameters : (i) Overall Perception and (ii) Country Perception. We asked the participants to identify the difference between the two parameters.

"<Overall sentiment> takes into just the entire tone of the paragraph where the country perception is specifically the tone towards that perceived country."-P6

"I kind of thought like country perception is like, I'm trying to look for what people think of a country as opposed to like the content of what happens in a certain country you are like" - P3

Participants indicated similar perceptions of the definitions. They were mostly in line with the ones given to them.

"Sometimes it comes off as being neutral or positive, but there must have been something negative about the country mentioned, so that was distinguishing that I made <between overall and country perception>. The article itself might have been wrapped up very positively." - P8

5.4.3 Diagnosis Parameter vs Toxicity Parameter. While the definitions of both parameters used very different adjectives to define them, the participants reduced them to higher and lower thresholds. This was the intention of the study to understand the different degrees of toxicity.

"But I think what I did was for diagnosis. It was like something little, not okay. For toxic, it was like, okay, this is problematic." - P8

6 DISCUSSION

6.1 Quantified Human Perception of Bias

Our results show that the annotators were able to identify the negativity and toxicity in the GPT-2 generated texts, even without knowing the source. Our Country and Overall Accentuation metrics show that the GPT-2 generated texts for the countries from Group-N showed a significant difference from the rest of the articles. They were perceived to be more negative and toxic than their human-written counterparts. Our adjective analysis also shows that GPT-2

has a perception of the Group-N countries that do not agree with its human-written counterpart, which shows an equal representation of positive and negative countries. GPT-2 written texts for Group-N countries heavily exaggerate military and war-like themes. We observed that our annotators could recall sharper details about these themes than the positive ones. This is indicative that negative texts have a more substantial impact on human memory. It also relates to the notion of implicit memory, which previous studies have indicated that individuals tend to prioritize the recollection of negative stereotypes over positive ones [5, 29, 53], resulting in an implicit inclination toward negative bias in memory retention. Consequently, our findings show the societal consequences of GPT-2 amplifying negative biases, as the false negative bias can lead readers to capture erroneous information.

Our work shows that the biases depicted by models like GPT-2 may have social impacts that also translate to representational and allotted harm. These results show that if such text generation models are used in a sociotechnical system, the biases identified can also be translated to potential harm. From the framework of harm postulated by Dev et al. [17], we can see how such behaviors can lead to harmful social behaviors such as *stereotyping*, *disparagement* and *erasure* where certain nationalities are oversimplified, evaluated as 'lesser' and underrepresented by the model respectively.

6.2 The Impressions of the Texts

Our qualitative results add explainability to our findings in the quantitative section. We use our qualitative results to dive deeper into 'why' the annotators answered the way they did. The results of our qualitative analysis and interview sessions reveal that the biases identified by the annotators did indeed create an impression of the country they were annotating, which in turn influenced their annotations moving forward. This highlights the instant impact of biases and their potential to shape how we view the world. This underscores the immediate influence of biases and their capacity to shape our perception of the world.

Interestingly, none of the annotators could explicitly identify the skewed perception of the country or that they were reading articles written by AI models until prompted to do so. Our analysis indicates that annotators were only able to implicitly identify bias by measuring the text as negative or toxic. This phenomenon can pose an issue as the biases in AI-generated content can remain unnoticed and continue to influence a reader's perceptions.

Another critical finding of our analysis was the phenomenon of 'AI hallucination', where AI models provide confident responses that seem faithful but are nonsensical in light of common knowledge [3]. A number of participants (prior to being prompted about the text source) mentioned that they felt that some of the passages were hard to follow and did not make any sense to them. They often reported that while the text began talking about one topic, the next topic would not be in line. Our study indicates that text generated by AI models tends to be influenced by AI hallucination, leading to a more radical and opinionated tone. This behavior makes AI-generated content more likely to mislead, as it is written in a confident and authoritative tone that can be perceived as factual by the reader [3, 49].

We also notice that the language style used by AI models correlates with how the model views a country. Specifically, GPT-2 generated stories about Group-N countries in a manner reminiscent of news articles, while it tends to present Group-P countries with a tone resembling opinion texts. This finding helps us understand how the model perceives the country and the associated information. An opinion piece by its very nature is perceived by the public as someone's opinion and not the ground truth. However, when N-Group country information is represented as news article it appears to be the ground truth. GPT-2 generated text further used terms like 'the BBC' to validate the idea being conveyed. This further leads to propagation of the AI authority phenomenon. This can have far reaching impacts when used in social scenarios. This finding underscores the importance of considering language style and its potential impact on perception when working with AI-generated content.

6.3 Human Perceptions of Automatic Indicators

Our qualitative analysis highlights the need for an interdisciplinary approach to bias identification in AI and NLP models. While previous studies have primarily focused on using automated evaluation to measure and quantify bias [10, 38], our paper presents a unique perspective by using human annotators to identify and attempt to quantify bias. In the process of conducting the interviews, we found that every participant had a different and unique perception of the goals of the study and the metrics they were asked to calculate. Although we use automatic indicators, our findings reveal that humans have differing perceptions on the given definitions of these parameters. These perceptions that do not always match those provided in the field of AI. Our readers, who could identify and rate bias in sentiment, viewed sentiment and differently, as seen by the low Cohen-Kappa values (OveP = 0.34, DiaP = 0.38). Similar values were observed for toxicity as well (DiaP = 0.18, ToxP = 0.38). These results show that it is necessary to consider differing human interpretations when defining and understanding computation-based parameters that can have a direct impact on human perceptions.

7 CONCLUSION

The paper uses human evaluation to explore nationality bias in a text generation model (GPT-2). The research uses the NLP sentiment and toxicity framework through human annotators to quantitatively analyze the presence of nationality bias. The findings reveal that the text generation model accentuates negative bias towards certain countries while demonstrating positive bias toward 'well-represented' countries. Using interviews, the study investigates how readers interpret articles generated by GPT-2. The interviews show that negative stories generated about certain countries had the most emotional impact on readers. However, some readers found such articles informative, informing them of a new aspect of the country. The study also found that participants were more welcoming of such technology after the disclosure that the articles were generated by both human and AI agents, as they were intended to 'mimic human behavior and biases.' The paper highlights the harmful impact of such technology if not used appropriately. It can enhance a country's skewed perception while maintaining the majority's viewpoint, leading to misinformation and stereotyping.

REFERENCES

- [1] Abubakar Abid, Maheen Farooqi, and James Zou. 2021. Persistent anti-muslim bias in large language models. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 298–306.
- [2] Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2019. Discrimination through optimization: How Facebook's Ad delivery can lead to biased outcomes. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–30.
- [3] Hussam Alkaiissi and Samy I McFarlane. 2023. Artificial Hallucinations in ChatGPT: Implications in Scientific Writing. *Cureus* 15, 2 (2023).
- [4] Joseph E Aoun. 2018. Optimism and Anxiety: Views on the Impact of Artificial Intelligence and Higher Education's Response. *Gallup Inc. Vol* (2018).
- [5] Mahzarin R Banaji and Anthony G Greenwald. 2016. *Blindspot: Hidden biases of good people*. Bantam.
- [6] Julia Barnett and Nicholas Diakopoulos. 2022. Crowdsourcing Impacts: Exploring the Utility of Crowds for Anticipating Societal Impacts of Algorithmic Decision Making. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*. 56–67.
- [7] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 610–623.
- [8] Lindsay Blackwell, Jill Dimond, Sarita Schoenebeck, and Cliff Lampe. 2017. Classification and its consequences for online harassment: Design insights from heartmob. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 1–19.
- [9] Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (Technology) is Power: A Critical Survey of "Bias" in NLP. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 5454–5476.
- [10] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems* 29 (2016).
- [11] Daniel Borkan, Lucas Dixon, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2019. Nuanced metrics for measuring unintended bias with real data for text classification. In *Companion proceedings of the 2019 world wide web conference*. 491–500.
- [12] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- [13] Aylin Caliskan, Pimparkar Parth Ajay, Tessa Charlesworth, Robert Wolfe, and Mahzarin R Banaji. 2022. Gender Bias in Word Embeddings: A Comprehensive Analysis of Frequency, Syntax, and Semantics. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*. 156–170.
- [14] Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356, 6334 (2017), 183–186.
- [15] Le Chen, Ruijun Ma, Anikó Hannák, and Christo Wilson. 2018. Investigating the impact of gender on rank in resume search engines. In *Proceedings of the 2018 chi conference on human factors in computing systems*. 1–14.
- [16] Mark Davies. 2017. The new 4.3 billion word NOW corpus, with 4–5 million words of data added every day. In *The 9th International Corpus Linguistics Conference*.
- [17] Sunipa Dev, Emily Sheng, Jieyu Zhao, Aubrie Amstutz, Jiao Sun, Yu Hou, Mattie Sanseverino, Jiin Kim, Akihiro Nishi, Nanyun Peng, et al. 2021. On Measures of Biases and Harms in NLP. *arXiv preprint arXiv:2108.03362* (2021).
- [18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. 4171–4186.
- [19] Quan Do. 2019. Jigsaw Unintended Bias in Toxicity Classification. (2019).
- [20] Jade S Franklin, Karan Bhanot, Mohamed Ghalwash, Kristin P Bennett, Jamie McCusker, and Deborah L McGuinness. 2022. An Ontology for Fairness Metrics. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*. 265–275.
- [21] Batya Friedman and Helen Nissenbaum. 1996. Bias in computer systems. *ACM Transactions on Information Systems (TOIS)* 14, 3 (1996), 330–347.
- [22] Ismael Garrido-Muñoz, Arturo Montejo-Ráez, Fernando Martínez-Santiago, and Alfonso Ureña-López. 2021. A survey on bias in deep nlp. *Applied Sciences* 11, 7 (2021), 3184.
- [23] Bhanu Prakash Reddy Guda, Aparna Garimella, and Niyati Chhaya. 2021. Empathbert: A bert-based framework for demographic-aware empathy prediction. *arXiv preprint arXiv:2102.00272* (2021).
- [24] Shloak Gupta, S Bolden, Jay Kachhadia, A Korsunskaya, and J Stromer-Galley. 2020. PoliBERT: Classifying political social media messages with BERT. In *Social, cultural and behavioral modeling (SBP-BRIMS 2020) conference*. Washington, DC.
- [25] Vipul Gupta, Pranav Narayanan Venkit, Shomir Wilson, and Rebecca J Passon-arau. 2023. Survey on Sociodemographic Bias in Natural Language Processing. *arXiv preprint arXiv:2306.08158* (2023).
- [26] Nida Manzoor Hakak, Mohsin Mohd, Mahira Kirmani, and Mudasir Mohd. 2017. Emotion analysis: A survey. In *2017 international conference on computer, communications and electronics (COMPELIX)*. IEEE, 397–402.
- [27] Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M Branham. 2018. Gender recognition or gender reductionism? The social implications of embedded gender recognition systems. In *Proceedings of the 2018 chi conference on human factors in computing systems*. 1–13.
- [28] Anikó Hannák, Claudia Wagner, David Garcia, Alan Mislove, Markus Strohmaier, and Christo Wilson. 2017. Bias in online freelance marketplaces: Evidence from taskrabbit and fiverr. In *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*. 1914–1933.
- [29] Richard L Hense, Louis A Penner, and Douglas L Nelson. 1995. Implicit memory for age stereotypes. *Social Cognition* 13, 4 (1995), 399–415.
- [30] Dirk Hovy. 2015. Demographic factors improve classification performance. In *Proceedings of the 53rd annual meeting of the Association for Computational Linguistics and the 7th international joint conference on natural language processing (volume 1: Long papers)*. 752–762.
- [31] Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denuyl. 2020. Social Biases in NLP Models as Barriers for Persons with Disabilities. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 5491–5501.
- [32] Clayton Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media*, Vol. 8. 216–225.
- [33] Lilly C Irani and M Six Silberman. 2013. Turkopticon: Interrupting worker invisibility in amazon mechanical turk. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 611–620.
- [34] Shivani Kapania, Oliver Siy, Gabe Clapper, Azhagu Meena SP, and Nithya Sambasivan. 2022. "Because AI is 100% right and safe": User Attitudes and Sources of AI Authority in India. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [35] Matthew Kay, Cynthia Matuszek, and Sean A Munson. 2015. Unequal representation and gender stereotypes in image search results for occupations. In *Proceedings of the 33rd annual acm conference on human factors in computing systems*. 3819–3828.
- [36] Patrick Gage Kelley, Yongwei Yang, Courtney Heldreth, Christopher Moessner, Aaron Sedley, Andreas Kramm, David T Newman, and Allison Woodruff. 2021. Exciting, useful, worrying, futuristic: Public perception of artificial intelligence in 8 countries. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 627–637.
- [37] Brendan Kennedy, Xisen Jin, Aida Mostafazadeh Davani, Morteza Dehghani, and Xiang Ren. 2020. Contextualizing Hate Speech Classifiers with Post-hoc Explanation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 5435–5442.
- [38] Svetlana Kiritchenko and Saif M Mohammad. 2018. Examining gender and race bias in two hundred sentiment analysis systems. *arXiv preprint arXiv:1805.04508* (2018).
- [39] Sarah Kreps, R Miles McCain, and Miles Brundage. 2022. All the news that's fit to fabricate: AI-generated text as a tool of media misinformation. *Journal of experimental political science* 9, 1 (2022), 104–117.
- [40] Juhi Kulkshrestha, Motahare Eslami, Johnatan Messias, Muhammad Bilal Zafar, Saptarshi Ghosh, Krishna P Gummadi, and Karrie Karahalios. 2017. Quantifying search bias: Investigating sources of bias for political searches in social media. In *Proceedings of the 2017 acm conference on computer supported cooperative work and social computing*. 417–432.
- [41] Keita Kurita, Nidhi Vyay, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. 2019. Measuring Bias in Contextualized Word Representations. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*. 166–172.
- [42] Bing Liu et al. 2010. Sentiment analysis and subjectivity. *Handbook of natural language processing* 2, 2010 (2010), 627–666.
- [43] Steven Loria. 2018. textblob Documentation. *Release 0.15.2* (2018).
- [44] Arm Ltd. 2017. AI Today, AI Tomorrow. Awareness and Anticipation of AI: A Global Perspective. <https://www.arm.com/solutions/artificial-intelligence/survey>
- [45] Luwei Rose Luqui and Fan Yang. 2018. Islamophobia in China: news coverage, stereotypes, and Chinese Muslims' perceptions of themselves and Islam. *Asian Journal of Communication* 28, 6 (2018), 598–619.
- [46] Milagros Miceli, Julian Posada, and Tianling Yang. 2022. Studying up machine learning data: Why talk about bias when we mean power? *Proceedings of the ACM on Human-Computer Interaction* 6, GROUP (2022), 1–14.
- [47] Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. StereoSet: Measuring stereotypical bias in pretrained language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 5356–5371.

- [48] Lisa-Maria Neudert, Aleksu Knuutila, and Philip N Howard. 2020. *Global attitudes towards AI, machine learning & automated decision making*. Technical Report. Working paper 2020.10, Oxford Commission on AI & Good Governance.
- [49] Andrew Ng. 2023. The Batch: ChatGPT Mania!, Crypto Fiasco Defunds AI Safety, Alexa Tells Bedtime Stories. <https://www.deeplearning.ai/the-batch/issue-174/>
- [50] Cathy O'neil. 2017. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.
- [51] Nedjma Ousidhoum, Xinran Zhao, Tianqing Fang, Yangqiu Song, and Dit-Yan Yeung. 2021. Probing Toxic Content in Large Pre-Trained Language Models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 4262–4274.
- [52] Bo Pang, Lillian Lee, et al. 2008. Opinion mining and sentiment analysis. *Foundations and Trends® in information retrieval* 2, 1–2 (2008), 1–135.
- [53] Charles W Perdue and Michael B Gurtman. 1990. Evidence for the automaticity of ageism. *Journal of Experimental Social Psychology* 26, 3 (1990), 199–216.
- [54] Vinodkumar Prabhakaran, Ben Hutchinson, and Margaret Mitchell. 2019. Perturbation Sensitivity Analysis to Detect Unintended Model Biases. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 5740–5745.
- [55] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.
- [56] William S Robinson. 1951. The logical structure of analytic induction. *American Sociological Review* 16, 6 (1951), 812–818.
- [57] Joel Ross, Lilly Irani, M Six Silberman, Andrew Zaldivar, and Bill Tomlinson. 2010. Who are the crowdworkers? Shifting demographics in Mechanical Turk. In *CHI'10 extended abstracts on Human factors in computing systems*. 2863–2872.
- [58] Niloufar Salehi, Lilly C Irani, Michael S Bernstein, Ali Alkhatib, Eva Ogbé, and Kristy Milland. 2015. We are dynamo: Overcoming stalling and friction in collective action for crowd workers. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 1621–1630.
- [59] Salim Sazzed. 2021. A Hybrid Approach of Opinion Mining and Comparative Linguistic Analysis of Restaurant Reviews. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*. 1281–1288.
- [60] Aaron Smith. 2018. Public attitudes toward computer algorithms. (2018).
- [61] Anna Sotnikova, Yang Trista Cao, Hal Daumé III, and Rachel Rudinger. 2021. Analyzing stereotypes in generative text inference tasks. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. 4052–4065.
- [62] Samuel J Stratton. 2021. Population research: convenience sampling strategies. *Prehospital and disaster Medicine* 36, 4 (2021), 373–374.
- [63] Pranav Narayanan Venkit, Sanjana Gautam, Ruchi Panchanadikar, Ting-Hao Huang, and Shomir Wilson. 2023. Nationality Bias in Text Generation. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*. 116–122.
- [64] Pranav Narayanan Venkit, Mukund Srinath, and Shomir Wilson. 2022. A Study of Implicit Bias in Pretrained Language Models against People with Disabilities. In *Proceedings of the 29th International Conference on Computational Linguistics*. 1324–1332.
- [65] Pranav Narayanan Venkit and Shomir Wilson. 2021. Identification of bias against people with disabilities in sentiment analysis and toxicity detection models. *arXiv preprint arXiv:2111.13259* (2021).
- [66] Mayur Wankhade, Annavarapu Chandra Sekhara Rao, and Chaitanya Kulkarni. 2022. A survey on sentiment analysis methods, applications, and challenges. *Artificial Intelligence Review* 55, 7 (2022), 5731–5780.
- [67] Meredith Whittaker, Meryl Alper, Cynthia L Bennett, Sara Hendren, Liz Kaziunas, Mara Mills, Meredith Ringel Morris, Joy Rankin, Emily Rogers, Marcel Salas, et al. 2019. Disability, bias, and AI. *AI Now Institute* (2019).
- [68] WorldBank. 2015. Individuals using the internet (% of population) - united states. <https://data.worldbank.org/indicator/IT.NET.USER.ZS?end=2017&locations=US&start=2015>
- [69] Florian Znaniecki. 1934. *The method of sociology*. Farrar & Rinehart.